

IDENTIFIABILITY OF 3-CLASS JUKES-CANTOR MIXTURES

COLBY LONG AND SETH SULLIVANT

ABSTRACT. We prove identifiability of the tree parameters of the 3-class Jukes-Cantor mixture model. The proof uses ideas from algebraic statistics, in particular: finding phylogenetic invariants that separate the varieties associated to different triples of trees; computing dimensions of the resulting phylogenetic varieties; and using the disentangling number to reduce to trees with a small number of leaves. Symbolic computation also plays a key role in handling the many different cases and finding relevant phylogenetic invariants.

1. INTRODUCTION

A phylogenetic model is a statistical model of the evolutionary relationships among a group of taxa. A standard feature of these models is a tree parameter which is meant to encode the common ancestry of the taxa under consideration. If a model produces a probability distribution consistent with observed data, one would like to be able to infer the true phylogeny from the probability distribution. In order for this to be possible, a distribution arising from the model must uniquely determine the tree parameter that produced it. In other words, one would like to be able to determine if the tree parameter of the model is *identifiable*.

The identifiability of both the tree parameter and the continuous parameters has already been established for the basic models of character evolution [7] as well as some of the more complex models [2, 3, 5]. The more complex phylogenetic models incorporate specific information about the mechanisms of evolution. For example, mixture models are designed to account for biological phenomena that result in data from more than one tree. A number of papers have examined the identifiability of mixture models with various restrictions on the topology of the tree parameters [6, 12, 13, 15]. Recent work on two-class mixture models has established the identifiability of the tree parameters for 2-class mixture of both the Jukes-Cantor and Kimura 2-parameter models [4]. Our goal in the present paper is to extend the ideas from [4] to larger class mixture models, in particular, we extend the results for 3-tree Jukes-Cantor mixtures. Our main result is the following:

Theorem 1.1. *The tree parameters of the 3-class Jukes-Cantor mixture model are identifiable on trees with ≥ 6 leaves.*

The proof of this main result will occupy the whole of the present paper and uses tools from algebraic geometry and combinatorics as well as some heavy symbolic computation.

In Section 2, we will demonstrate why algebraic geometry is the appropriate tool for studying these models by associating to each set of tree parameters an algebraic variety of the possible distributions arising from those trees. We will then show how the question of identifiability can be reduced to showing that for any two sets of tree parameters, the associated varieties are not contained in one another. To show the varieties are not contained in one another, it is enough to show that their vanishing ideals are not contained in one another. Elements of the vanishing ideal are called *phylogenetic invariants*, and isolating these invariants will be a key part of this proof.

In Section 3 we will investigate the combinatorial properties of binary leaf-labelled trees to show that it is not actually necessary to compare *arbitrary* sets of tree parameters. Instead, we will be able to obtain general identifiability results for n -leaf trees by comparing mixtures on

trees with ≤ 6 leaves. Thus, we will have a finite list of pairs of mixtures for which we must show the mutual noncontainment of their varieties.

In Section 4 we will introduce the Fourier-Hadamard transformation which will simplify the parameterization of our varieties by turning them into toric varieties. We will use these new coordinates in Section 5 to show that for n -leaf trees, all of the varieties under consideration have the same dimension. This observation halves the number of phylogenetic invariants we must find to separate two mixtures.

Finally, in Section 6, we will combine the results from the previous sections to construct a finite list of specific pairs of mixtures that we must consider. We will then outline a method for finding phylogenetic invariants that distinguish these mixtures from one another and provide access to computations proving that they exist. Many, but not all pairs of triplets of trees are separated by linear invariants. For the triplets not separated by linear invariants, we will use the linear invariants in a novel way to construct separating invariants of higher degree.

2. PRELIMINARIES

The tree parameter of a phylogenetic model of k -state character change on an n -taxon tree is a binary leaf-labelled tree T with label set $[n]$. To each node v of the tree, we associate a random variable X_v that can take on any of the k characters. For the unique edge e between v and w , the matrix A^e is the matrix with (i, j) -th entry equal to $P(X_v = i | X_w = j)$. In other words, this matrix encodes the probability of a character transition along the edge e . Note that there are k^n possible character states at the leaves, and we can calculate the probability of observing a particular state by summing over all possible states of the internal nodes. The continuous parameters of such a model are the entries of these matrices, and for each choice of parameters we obtain a probability distribution at the leaves. Therefore, we have a map from the stochastic parameters Θ_T into the probability simplex Δ^{k^n-1} ,

$$\psi_T : \Theta_T \rightarrow \Delta^{k^n-1} \subseteq \mathbb{R}^{k^n}.$$

Since every element of $\text{im}(\psi_T)$ is a probability distribution, all of the entries must be between zero and one and must sum to one. Each coordinate function of ψ_T is a polynomial map in the continuous parameters. The degree of each polynomial is equal to the number of edges in the tree parameter and the number of terms is equal to k raised to the number of internal vertices. Therefore, if we ignore the restrictions on the domain and range and simply regard ψ_T as a complex polynomial map, $\overline{\text{im}(\psi_T)} = V_T$ is an algebraic variety.

An r -class mixture model enlarges the space of possible distributions by taking r tree parameters as input and introducing $r - 1$ *mixing parameters*. The mixing parameters weight the distribution from each of the trees according to the proportion of data arising from that tree. Note that the underlying trees might not be distinct. This could be the case where the tree topologies are the same, but the entries of the transition matrices are not the same along corresponding edges. Just as before, with fixed tree parameters, we have a map that takes a choice of continuous parameters for each tree and a choice of mixing parameters and sends them to a probability distribution. As an example, consider the map for a 3-class mixture, which will be our primary object of interest.

$$\psi_{T_1, T_2, T_3} : \Theta_{T_1} \times \Theta_{T_2} \times \Theta_{T_3} \times \Delta^2 \rightarrow \Delta^{k^n-1}$$

where

$$(s_1, s_2, s_3, \pi) \mapsto \pi_1 \psi_1(s_1) + \pi_2 \psi_2(s_2) + \pi_3 \psi_3(s_3).$$

Here, $\pi = (\pi_1, \pi_2, \pi_3) \in \Delta^2$ is the vector of mixing parameters. Again, regarded as a complex polynomial map, $\text{im}(\psi_{T_1, T_2, T_3})$ is an algebraic variety. In fact,

$$\overline{\text{im}(\psi_{T_1, T_2, T_3})} = V_{T_1} * V_{T_2} * V_{T_3},$$

where $V_{T_1} * V_{T_2} * V_{T_3}$ denotes the *join* variety of V_{T_1} , V_{T_2} , and V_{T_3} .

Before we formally define the concept of identifiability for r -class mixtures, we will introduce some convenient notation. Let \mathcal{T}_X be the set of trivalent leaf-labelled trees with label set X and let $\mathcal{T}_{X,r}$ be the set of unordered lists of elements of \mathcal{T}_X of length r (i.e. r element multisets). Note that as in our mixture models, for $T = (T_1, \dots, T_r)$, the trees in T are not necessarily distinct. We will now write $\psi_T := \psi_{T_1, \dots, T_r}$.

Definition 2.1. [4] The tree parameters of an r -tree mixture model are *generically identifiable* for n leaf trees if for all $T \in \mathcal{T}_{[n],r}$ and generic choices of $(s_1, \dots, s_r, \pi) \in \Theta_{T_1} \times \dots \times \Theta_{T_r} \times \Delta^{r-1}$, if there is a T' and $(s'_1, \dots, s'_r, \pi') \in \Theta_{T'_1} \times \dots \times \Theta_{T'_r} \times \Delta^{r-1}$ such that

$$\psi_T(s_1, \dots, s_r, \pi) = \psi_{T'}(s'_1, \dots, s'_r, \pi')$$

then $T = T'$.

In this paper we will not need such generality, as we will specifically consider the 3-class Jukes-Cantor mixture model. The Jukes-Cantor model is a 4-state character substitution model of DNA evolution, with states A, C, G, and T corresponding to the DNA bases. We assume equal transition probabilities between characters, so the transition matrices have the form

$$\begin{pmatrix} \alpha & \beta & \beta & \beta \\ \beta & \alpha & \beta & \beta \\ \beta & \beta & \alpha & \beta \\ \beta & \beta & \beta & \alpha \end{pmatrix}.$$

Because the entries of each row of the transition matrix must sum to one, we essentially have one parameter along each edge, though we will often ignore this in order to homogenize the parameterization. In this context, we can think of the parameter value on an edge as representing edge length, with greater β values corresponding to a higher probability of mutation and longer branches.

In order to prove Theorem 1.1, we will translate this statement about identifiability into one about algebraic varieties.

Lemma 2.2. [4] Suppose T_1, T_2, T_3, S_1, S_2 , and S_3 are binary, n -leaf trees, not necessarily distinct, then for the 3-tree Jukes-Cantor mixture model, $V_{T_1} * V_{T_2} * V_{T_3} \not\subseteq V_{S_1} * V_{S_2} * V_{S_3}$ and $V_{S_1} * V_{S_2} * V_{S_3} \not\subseteq V_{T_1} * V_{T_2} * V_{T_3}$ implies that the set of stochastic parameters mapping into $V_{S_1} * V_{S_2} * V_{S_3} \cap V_{T_1} * V_{T_2} * V_{T_3}$ is a set of Lebesgue measure 0.

This algebraic characterization means that we are able to obtain results about the stochastic parameters by working with complex varieties and all the tools thereof. One strategy for proving generic identifiability of the 3-class Jukes-Cantor mixture model on n -taxa is then clear. We can simply list all elements of $\mathcal{T}_{[n],3}$ (which we will call *n -leaf triplets*) and generate the corresponding varieties. By Lemma 2.2, if we can show that any two of these varieties are mutually noncontained, then we will have established identifiability in the n -leaf case. As alluded to in the introduction, we will actually want to look at elements of $I(V_{T_1} * V_{T_2} * V_{T_3})$, which we call the phylogenetic invariants of $V_{T_1} * V_{T_2} * V_{T_3}$, (or occasionally just phylogenetic invariants of $T = \{T_1, T_2, T_3\}$, or of the mixture model). Therefore, for each $(S, T) \in \mathcal{T}_{[n],r} \times \mathcal{T}_{[n],r}$ with $S \neq T$, we need to find an invariant of T that is not an invariant of S , and vice versa. Once we have done this for a specific pair, we will say that we have *separated* S and T .

This gives us a clear procedure for determining identifiability, but with some obvious drawbacks. First, the number of binary phylogenetic trees on n -taxa is $(2n-5)!! = 1 \cdot 3 \cdot 5 \cdots (2n-5)$, which makes generating all possible 3-tree mixtures computationally prohibitive even for relatively small n . Secondly, on the face of it, this brute force approach does not seem to offer any way of establishing generic identifiability of the model for arbitrary n . However, as we will see

in the next section, it is possible to establish generic identifiability of the 3-tree Jukes-Cantor mixture model for all n by separating only a finite number of mixtures.

3. DISENTANGLING 3-TREE MIXTURES

In this section we explain how to use trees with few leaves to establish identifiability. The size of the trees we need to consider is bounded by the disentangling number for phylogenetic mixtures. For $T \in \mathcal{T}_X$ and $K \subset X$, let $T|_K$ be the tree obtained by supressing all degree two vertices in the subtree of T induced by the leaves labelled by K . For $T = (T_1, \dots, T_r) \in \mathcal{T}_{X,r}$, $T|_K = (T_{1|K}, \dots, T_{r|K})$.

Example 3.1. Consider $T \in \mathcal{T}_{[8]}$ pictured below and $K = \{2, 3, 5, 7, 8\}$.



Definition 3.2. Let $S, T \in \mathcal{T}_{X,r}$ with $S \neq T$. A subset $K \subseteq X$ is said to disentangle S and T if $S|_K \neq T|_K$. Let $d(S, T)$ be the cardinality of the minimum disentangling set of S and T . The disentangling number $D(r)$ is

$$D(r) = \max_{n \in \mathbb{N}} \max_{S \neq T \in \mathcal{T}_{[n],r}} d(S, T)$$

The following lemma [4] motivates our interest in the disentangling number.

Lemma 3.3. Let $S, T \in \mathcal{T}_{[n],3}$ and $K \subseteq [n]$. If $V_{T_{1|K}} * V_{T_{2|K}} * V_{T_{3|K}} \not\subseteq V_{S_{1|K}} * V_{S_{2|K}} * V_{S_{3|K}}$ then $V_{T_1} * V_{T_2} * V_{T_3} \not\subseteq V_{S_1} * V_{S_2} * V_{S_3}$.

Now suppose we are able to show identifiability for trees with $D(3)$ leaves. Then given any two mixtures S, T on $n > D(3)$ leaves we can find some $K \subset [n]$ such that $T|_K \neq S|_K$ and $V_{T_{1|K}} * V_{T_{2|K}} * V_{T_{3|K}} \not\subseteq V_{S_{1|K}} * V_{S_{2|K}} * V_{S_{3|K}}$. By Lemma 3.3, in doing so we will have separated S and T . Consequently, we would have generic identifiability of the tree parameter of the 3-class Jukes-Cantor mixture model for all $n \geq D(3)$. Thus, as promised, we will have an upper bound on the number of possible varieties we need to consider. In this section we provide some general background on the disentangling number and prove that $D(3) = 6$.

The *rooted* disentangling number, $RD(r)$, is defined analogously for rooted trees. We will omit the short proof of this lemma from [19] that relates $RD(r)$ and $D(r)$.

Lemma 3.4. The disentangling and rooted disentangling numbers satisfy: $D(r) \leq RD(r) + 1$.



The main result of [19] is the following theorem from which we obtain an upper bound on $D(r)$ as an immediate corollary.

Theorem 3.5. $RD(r) = 3(\lfloor \log_2(r) \rfloor + 1)$.

Corollary 3.6. For $r \in \mathbb{N}$, $D(r) \leq 3(\lfloor \log_2(r) \rfloor + 1) + 1$.

The original proof Theorem 3.5 is obtained by encoding multisets of trees as high-dimensional contingency tables and applying results about marginal maps. We provide an alternative, and hopefully more direct proof by examining the tree topologies directly.

Proof of Theorem 3.5. A construction in [11] shows that $3(\lfloor \log_2(r) \rfloor + 1) \leq RD(r)$, so we need only show that for every pair $S, T \in \mathcal{T}_{X,r}$ there is a disentangling set of cardinality less than or equal to $3(\lfloor \log_2(r) \rfloor + 1)$. We will proceed by induction on r . Because a rooted tree is determined

FIGURE 1. Possible locations for e_1 .FIGURE 2. Possible structures for 3-partners at $\{K_1, K_2, K_3\}$.

by its rooted triples ([16, Theorem 6.4.1]), the base case $RD(1) = 3$ is established. Assume this is true for all integers less than r and let $S = (S_1, \dots, S_r)$ and $T = (T_1, \dots, T_r)$ be two unordered lists of rooted binary trees with $S \neq T$. There must exist some S_i and T_j such that $S_i \neq T_j$. By our inductive assumption, we can permute the leaf labels so that for $K = \{1, 2, 3\}$, $T_{i|K} \neq S_{j|K}$. There are only three topologically distinct rooted leaf-labelled binary trees with label set K , which we will label t_1, t_2 , and t_3 .

If $S_{|K} \neq T_{|K}$, then S and T are disentangled and we are done. Otherwise, $S_{|K} = T_{|K}$ is an unordered list of the trees t_1, t_2 , and t_3 occurring with multiplicity. Partition S into three multisets,

$$L_S^l := \{S_j \in S \mid S_{j|K} = t_l\}$$

for $1 \leq l \leq 3$, and likewise for T . Since K was chosen to disentangle an element of S from an element of T , it must be the case that $S_{|K}$ and $T_{|K}$ contain at least two distinct three leaf trees. Therefore, we can choose l so that L_S^l is nonempty and $|L_S^l| = r' \leq \frac{r}{2}$. Since $S_{|K} \neq T_{|K}$, $|L_S^l| = |L_T^l|$ and we can consider L_T^l and L_S^l as elements of $\mathcal{T}_{X, r'}$. By our inductive assumption, there exists a disentangling set K' such that

$$\begin{aligned} |K'| &\leq 3(\lfloor \log_2(r') \rfloor + 1) \\ &\leq 3(\lfloor \log_2(\frac{r}{2}) \rfloor + 1) \\ &= 3((\lfloor \log_2(r) \rfloor - 1) + 1) \\ &= 3(\lfloor \log_2(r) \rfloor). \end{aligned}$$

Therefore, $|K \cup K'| \leq 3(\lfloor \log_2(r) \rfloor + 1)$. We claim that this set disentangles S and T . Since K' disentangles L_T^l from L_S^l , and $K' \subseteq K \cup K'$, $(L_S^l)_{|K \cup K'} \neq (L_T^l)_{|K \cup K'}$. If S and T are still entangled, then there must be some tree in $(L_S^l)_{|K \cup K'}$ equal to some tree in $(L_T^m)_{|K \cup K'}$ with $l \neq m$. But since $K \subseteq K \cup K'$, this is impossible, so $K \cup K'$ disentangles S and T . \square

While this assures us that $D(3) \leq 7$, we can actually reduce this bound slightly, vastly reducing the number of 3-tree mixtures we need to consider. For the theorem and proof that follow, we will make use of the following definition.

Definition 3.7. For $T \in \mathcal{T}_{[n]}$ and K a three element subset of $[n]$, if T has a split that separates K from $[n] \setminus K$ then $T_{|K}$ is a *cluster on K* .

Definition 3.8. Let $S, T \in \mathcal{T}_{X, r}$ and K a subset of X that does not disentangle S and T . Label the trees of S and T so that $S = (S_1, \dots, S_r)$ and $T = (T_1, \dots, T_r)$. For $1 \leq m \leq r$, let

$$m_i = \min(\{m \in [r] \setminus \{m_1, \dots, m_{i-1}\} : S_{m|K} = T_{i|K}\})$$

Then with respect to the chosen labeling, we say that S_{m_i} and T_i are *partners* at K .

Notice that each tree of T has exactly one partner at K , and that the partnered trees at K are exactly the same if we swap the roles of S and T in the definition.

Theorem 3.9. $D(3) \leq 6$.




FIGURE 3. Structure of trees satisfying Case 1.




FIGURE 4. Structure of trees satisfying Case 2.

Proof. We will use contradiction. Suppose $D(r) = 7$ and let $K_i = [7] \setminus i$. Then there must exist $S, T \in \mathcal{T}_{[7],3}$ such that $S|_{K_i} = T|_{K_i}$ for $1 \leq i \leq 7$. For everything that follows, fix some labelling of the trees of S and T so that for each i , every tree of T and S has a partner at K_i .

We will collect a few key observations about trees that are partnered together at multiple K_i . If a tree of S and a tree of T are partnered together at K_i for exactly j distinct values of i , then we will call them j -partners. Suppose $S_l \neq T_m$ are 2-partners and permute the leaf labels so that they are partnered at K_1 and at K_2 . Let v_i be the leaf vertex labelled i , and let e_i be the edge adjacent to this vertex. Since $S_l|_{K_1} = T_m|_{K_1}$, there must be a unique edge on this tree where e_1 is attached to form S_l , and a different unique edge where e_1 is attached to form T_m . But in order for our trees to still be equal when restricted to K_2 , the two distinct edges where we attached e_1 must collapse to the same edge when we remove e_2 . Therefore, S_l and T_m must have the structure of the tree in Figure 1 where r_1 and r_2 are rooted trees and where e_1 is one of the dashed edges.

Now suppose that S_l and T_m are 3-partners, partnered at K_1, K_2 and K_3 . From above, observe that regardless of which edge is e_1 , that the length of the path from v_1 to v_2 in both trees must be less than or equal to three. Therefore, the length of the path between each pair of vertices, (v_1, v_3) , (v_1, v_2) , and (v_2, v_3) , must be less than three in both S_l and T_m . Consequently, S_l and T_m must both have a cluster on $\{1, 2, 3\}$, and they must be the same tree apart from these clusters (i.e., S_l and T_m must be two different trees from the list in Figure 2, where r is some rooted tree). From the figure we also see if $S_l|_{K_i} = T_m|_{K_i}$ for $i \notin \{1, 2, 3\}$, then $S_l = T_m$. This implies that for $j > 3$ any two trees that are j -partners must be the same tree.

If any tree of S is equal to any tree of T , then we can remove these trees to form the lists $S', T' \in \mathcal{T}_{[7],2}$, and any set K that disentangles S' and T' will disentangle S and T . Since $D(2) = 6$ ([12]) this would imply $d(S, T) \leq 6$ contradicting our assumption that $d(S, T) = 7$. Therefore, we can assume that no two trees are j -partners for $j > 3$. Since each tree of S and each tree of T must be partnered at all seven K_i , the only possibility for a single tree is that it has one 3-partner and two 2-partners or two 3-partners and one 1-partner. The particular partnering relationships impose restrictions on the possible structures of the trees in S and T . We will now consider both cases and use these restrictions to arrive at a contradiction.

Case 1: There exists a tree in S or T with one 3-partner and two 2-partners.

We will leave all partners fixed according to the original labelling. However, for convenience, we will relabel the lists, trees, and leaves so that T_1 is partnered with S_1 at $\{K_1, K_2, K_3\}$, with S_2 at $\{K_4, K_5\}$, and with S_3 at $\{K_6, K_7\}$.

So far then, we know T_1 and S_1 are as in Figure 3, where t_1 and t_2 are distinct clusters on $\{1, 2, 3\}$ and r_1 is a rooted tree with leaf label set $\{4, 5, 6, 7\}$. We also know that $T_1|_{K_4} = S_2|_{K_4}$, so $S_2|_{K_4}$ contains the cluster t_1 . If e_4 is connected to an edge of S_2 somewhere in the cluster t_1

of $S_{2|K_4}$, then it is impossible for $S_{2|K_5} = T_1$. Therefore, we see that even without restricting to K_4 , S_2 and similarly S_3 must contain the cluster t_1 . Thus, T_1, S_1, S_2 , and S_3 are all as depicted in Figure 3.

Every tree must have a 3-partner, so let T_2 be a 3-partner of S_2 . From our observations above, S_2 and T_2 differ only by a cluster on some three element set K' . We know that K' cannot contain 4 or 5 since S_2 is partnered with T_1 at K_4 and K_5 . Therefore, K' contains at least one element of $\{1, 2, 3\}$. But to preserve the cluster t_1 , it must be that $K' = \{1, 2, 3\}$. If $T_{2|K'} = t_3$, then K' disentangles S and T , and if $T_{2|K'} = t_1$ then $T_2 = S_2$, it follows that $T_{2|K'} = t_2$.

Finally, S_1 and T_1 as well as S_2 and T_2 are partnered at K_1, K_2 , and K_3 , which forces S_3 and T_3 to be partnered at K_1, K_2 , and K_3 . As a result, S_3 and T_3 differ only by a cluster on $\{1, 2, 3\}$ (Figure 3). If $T_{3|K'} = t_1$ then $T_3 = S_3$ and if $T_3 = t_2$ or $T_3 = t_3$ then K' is a disentangling set. In any case we have a contradiction.

Case 2: Every tree in S and T has two 3-partners and one 1-partner.

As before, leave the partnering relationships fixed and relabel the lists, trees, and leaves so that T_1 is partnered with S_1 at $\{K_1, K_2, K_3\}$, with S_2 at $\{K_4, K_5, K_6\}$, and with S_3 at $\{K_7\}$. As we've seen, T_1 must have a cluster on $\{1, 2, 3\}$ and a cluster on $\{4, 5, 6\}$, so T_1 is as pictured in Figure 4. As 3-partners of T_1 , both S_1 and S_2 must have clusters on $\{1, 2, 3\}$ and $\{4, 5, 6\}$ as well and are also as depicted in Figure 4.

S_1 must have another 3-partner which we will label T_2 . S_1 and T_2 must differ only at a cluster on some three element set $K' \subset [7] \setminus \{1, 2, 3\}$ which must contain elements of $\{4, 5, 6\}$. Then as we argued above, to preserve the $\{4, 5, 6\}$ cluster on S_1 it must be that $K' = \{4, 5, 6\}$. S_2 must also have a second 3-partner. This tree can't be T_2 , since then both T_1 and T_2 would have two 3-partners, leaving T_3 as the sole 3-partner for S_3 and putting us back in Case 1. Therefore, S_2 and T_3 must be 3-partners, and the same logic shows that they differ only at a cluster on $\{1, 2, 3\}$. The possible structures of T_1, T_2, T_3, S_1 , and S_2 are all displayed in Figure 4 ($t_{i/j}$ indicates that a cluster can be only either t_i or t_j and likewise for $s_{i/j}$). Since the 3-partners of S_1 are T_1 and T_2 , the 1-partner of S_1 must be T_3 , and S_1 and T_3 must be partnered at K_7 . From the diagram, it is clear $S_{1|K_7} \neq T_{3|K_7}$, which is a contradiction. \square

4. THE FOURIER-HADAMARD COORDINATE TRANSFORMATION

For group-based models the Fourier-Hadamard coordinate transformation is a linear change of coordinates that makes each coordinate function of the parameterization a monomial. Importantly for our purposes, the linearity of the transformation means that it commutes with taking mixtures. In this description, both determining the dimension and finding phylogenetic invariants of the mixture varieties become much simpler. We will present a practical outline demonstrating how to recover the monomials; a thorough explanation of the transform can be found in [9], [18], and [20].

Let B and B' be the partition of $[n]$ induced by removing an edge e from a leaf-labelled tree T . The set of such splits of T is denoted $\Sigma(T)$ and uniquely determines the tree topology. As a result, we will index the edges of T by the splits that they induce. A phylogenetic model is group-based if there exists a group G and functions $f_{B|B'} : G \rightarrow \mathbb{R}$ associated to each edge of the n -leaf tree parameter T , such that when the character states are identified with the elements of G , the probability of character change along e is dependent only on the difference between the character states at the endpoints of e . In other words, the probability of the endpoints of e being in character states identified with group elements g and h is equal to $f_{B|B'}(g - h)$.

Let p_{g_1, \dots, g_n} be the probability of observing the state (g_1, \dots, g_n) at the leaves of T and let q_{g_1, \dots, g_n} be the image of this coordinate after the Fourier-Hadamard transformation. For the Jukes-Cantor model, we make the following identification of the character states with the

elements of $\mathbb{Z}_2 \times \mathbb{Z}_2$: $A \mapsto (0, 0)$, $C \mapsto (1, 0)$, $G \mapsto (0, 1)$, $T \mapsto (1, 1)$. Since the JC model assumes all transition probabilities are equal, our concern is only whether or not there are different characters at the endpoints of each edge. Thus, $f_{B|B'}((1, 0)) = f_{B|B'}((0, 1)) = f_{B|B'}((1, 1))$, and after transformation, the new coordinates of the parameter space are $a_C^{B|B'} = a_G^{B|B'} = a_T^{B|B'} \in (0, 1]$ and $a_A^{B|B'} = 1$. Then


$$q_{g_1, \dots, g_n} = \begin{cases} \prod_{B|B' \in \Sigma(T)} a_{\sum_{i \in B} g_i}^{B|B'} & : \sum_{i=1}^n g_i = 0 \\ 0 & : \text{otherwise} \end{cases}$$

Notice that in the group $\mathbb{Z}_2 \times \mathbb{Z}_2$, if the leaf elements sum to the identity then for every partition $B|B'$, $\sum_{i \in B} g_i = \sum_{i \in B'} g_i$. Therefore, the monomial above does not depend on our labelling of the sets of the splits. We call the splits of T associated to leaf edges the trivial splits and for simplicity we denote the parameters associated to the leaf edge labelled i by a_g^i .

Note that for each nontrivial coordinate and for each edge either $a_A^{B|B'}$ or $a_C^{B|B'}$ appears in the monomial parameterization of the coordinate. We encode the resulting monomials in tree diagrams as follows. Redraw the tree T , but make each edge solid if $a_C^{B|B'}$ appears and dotted if $a_A^{B|B'}$ appears. The solid edges of the diagram form a subforest of T and the number of distinct nontrivial Fourier coordinates are in bijection with the subforests of T [17].

Example 4.1. Let T be the tree with nontrivial splits given by $\{15|2346\}$, $\{135|246\}$, and $\{1235|46\}$. The parameterization of a particular coordinate as well as the subforest induced by this coordinate are shown.

$$q_{CAGGTG} = a_C^1 a_A^2 a_C^3 a_C^4 a_C^5 a_C^6 a_C^{15|2346} a_A^{135|246} a_A^{1235|46}$$



The inclusion of the $a_A^{B|B'}$ variables homogenizes the parameterization which will be convenient for some the work that follows. However, when writing out the explicit parameterization of a coordinate we will dehomogenize by setting all of the $a_A^{B|B'} = 1$. This allows us to drop the subscripts entirely. For arbitrary n -leaf trees, we will further simplify notation by extending the numbering of the trivial splits to a numbering of all of the splits so that every split is labelled by some element of the set $[2n - 3]$. Thus, in Example 4.1 the coordinate can be written more compactly as $q_{CAGGTG} = a_1 a_3 a_4 a_5 a_6 a_7$.

5. DIMENSION

Our goal in this section is to prove dimension results for the appropriate join varieties, which will allow us, for each pair of triplets, to prove non-containment in only one direction. Our technique for deducing the dimension results relies on the fact the the underlying varieties V_T are toric varieties as explained in the previous section. We make use of the tropical secant varieties approach of Draisma [8] to prove our dimension results and refer the reader to that paper for background.

When taking the join of r varieties, we introduce $(r - 1)$ parameters, giving us the following bound

$$\dim(V_1 * \dots * V_r) \leq \min \left\{ \sum_{i=1}^k \dim(V_i) + (r - 1), \dim(V) \right\},$$

where V is the ambient affine space containing the join variety. This bound is called the *expected dimension* and any join variety realizing this bound is called *nondefective*. As an example of how we will use this concept, let $T = (T_1, T_2, T_3)$ be a 6-leaf triplet. We know that each tree in T has 9 edges, and so from the Fourier parameterization of the Jukes-Cantor model, it is obvious that $\dim(V_{T_i}) \leq 9$. Therefore, if we can show that $\dim(V_{T_1} * V_{T_2} * V_{T_3}) = 29$, we will have shown that this variety is nondefective. Now suppose we could show that the join variety of every 6-leaf triplet is nondefective, then as a corollary, the join varieties of any two 6-leaf triplets are the same dimension. Since the varieties involved are irreducible, if they are of the same dimension, the containment

$$V_{T_1} * V_{T_2} * V_{T_3} \subseteq V_{S_1} * V_{S_2} * V_{S_3}$$

implies that the two join varieties are actually equal. Thus, to separate S and T , it is enough to show $V_{T_1} * V_{T_2} * V_{T_3}$ and $V_{S_1} * V_{S_2} * V_{S_3}$ are not the same variety, meaning we would only need to find an invariant for one that is not an invariant for the other. Since eventually our proof will necessitate separating certain triplet pairs, establishing nondefectiveness will greatly reduce the number of invariants we have to find. Moreover, it will make the task of separating pairs much easier in cases where finding an invariant to establish noncontainment in one direction is more difficult than doing so in the other.

The following definitions and terminology up to Theorem 5.3 are adapted from the more general presentation in [8]. Our approach is inspired by proof of the nondefectiveness of the second secant variety associated to a 4-leaf tree as shown in ([4]). Let $T \in \mathcal{T}_{[n],r}$, we will temporarily regard the a_A^i as variables which homogenizes the parameterization of V_{T_i} for $1 \leq i \leq r$. We will consider these projective varieties as affine cones, C_i , which are closed under scalar multiplication. Now the additional parameters introduced when constructing the join variety are superfluous, and we can take as our definition that the *join* of C_1 through C_r is

$$C_1 + \dots + C_r := \{c_1 + \dots + c_r \mid c_i \in C_i, 1 \leq i \leq r\}.$$

Each of these cones is the image of a polynomial map, $f_i : \mathbb{C}^{m_i} \rightarrow \mathbb{C}^q$, and in our case, the use of Fourier coordinates ensure that each coordinate function is a monomial. Then for $1 \leq i \leq r$, we can write f_i as a list $(f_{i,b})_{b=1}^q$ where each $f_{i,b} = x^\alpha$. For example, if we let $r = 3$ and $n = 6$, then $m_1 = m_2 = m_3 = 18$, $x = (a_C^1, \dots, a_C^9, a_A^1, \dots, a_A^9)$, $\alpha \in \{0, 1\}^{18}$, and $q = 4^6 = 4096$.

Let $l_{i,b} : \mathbb{R}^{m_i} \rightarrow \mathbb{R}$ be the piecewise linear function defined by

$$v \mapsto \langle v, \alpha \rangle.$$

For $v = (v_1, \dots, v_r) \in \prod_{i=1}^r \mathbb{R}^{m_i}$ and $b \in [q]$, we say that i *wins* b at v if

- (1) $l_{i,b}(v_i) < l_{j,b}(v_j)$ for all $j \neq i$ and
- (2) $l_{i,b}$ is differentiable (hence linear) near v_i .

Denote the corresponding differential by $d_{v_i} l_{i,b}$.

Definition 5.1. For $v = (v_1, \dots, v_r) \in \prod_{i=1}^r \mathbb{R}^{m_i}$, let

$$D_i(v) := \{d_{v_i} l_{i,b} \mid i \text{ wins } b \text{ at } v\}$$

be the set of *winning directions* of i at v .

Finally, we have all the requisite definitions to state the primary result we will need.

Lemma 5.2. ([8]) *The dimension of $C_1 + \dots + C_r$ is at least the maximum, taken over all $v = (v_1, \dots, v_r) \in \prod_{i=1}^r \mathbb{R}^{m_i}$, of the sum*

$$\sum_{i=1}^r \dim_{\mathbb{R}} \langle D_i(v) \rangle_{\mathbb{R}}.$$

This gives us a way to compute lower bounds on the dimensions of join varieties. Now to show that a join variety is non-defective, we just need to show that this lower bound is equal to the expected dimension.

Theorem 5.3. *Let $T \in \mathcal{T}_{[n],r}$. For $n \geq 4$ and $r \leq \lceil \frac{n}{2} \rceil$, the join variety $V_{T_1} * \dots * V_{T_r}$ associated to the r -class Jukes-Cantor mixture model is nondefective.*

Proof. Let $T \in \mathcal{T}_{[n],r}$, by Lemma 5.2, to show nondefectiveness it will be enough to find a vector $v = (v_1, \dots, v_r)$ so that for $1 \leq i \leq r$, $\dim_{\mathbb{R}} \langle D_i(v) \rangle_{\mathbb{R}} = (2n - 3) + 1$. Thus, $C_1 + \dots + C_r$ will have dimension $r(2n - 3) + r$, and when we projectivize by setting the $a_A^j = 1$, we will have $\dim(V_{T_1} * \dots * V_{T_r}) = r(2n - 3) + (r - 1)$ as desired.

The set of winning directions of i at v , $D_i(v)$, is a set of 0-1 vectors in \mathbb{R}^{4n-6} . Our goal will be to construct the vector v in such a way that the vectors in each $D_i(v)$ span a space of dimension $2n - 2$. Recall that for a tree T , the distinct Fourier coordinates are in bijection with the subforests of T . Therefore, each b induces a subforest on the trees T_1, \dots, T_r the number of leaf edges of which, t_b , depends only on the number of entries that are not A in the index of the b -th coordinate. For example, if the b -th coordinate is q_{AACGT} then $t_b = 3$ since in *any* 5-leaf tree the subforest induced by b contains 3 leaf edges.

Case 1: r is even

Construct the vector $v = (v_1, \dots, v_r)$ as follows. Each v_i has $4n - 6$ entries corresponding to leaf edges, half of which correspond to the variables a_C^j and half to the homogenizing variables a_A^j . Let the entries of v_i corresponding to leaf edges a_C^j be equal to α_i , to leaf edges a_A^j equal to β_i , and set all other entries equal to zero.

Then for $1 \leq j \leq r$, $1 \leq b \leq q$,

$$l_{i,b}(v_i) = \alpha_i t_b + (n - t_b) \beta_i = (\alpha_i - \beta_i) t_b + n \beta_i.$$

Notice that this function depends only on the number of leaf edges in the subforest of T_i induced by b . Let $\mu_i : \mathbb{R} \rightarrow \mathbb{R}$ be given by

$$t \mapsto (\alpha_i - \beta_i) t + n \beta_i.$$

The two parameters α_i and β_i allow us to make μ_i whatever line we wish in \mathbb{R}^2 . Now we have $l_{i,b}(v_i) = \mu_i(t_b)$, and as explained, this value is completely independent of the trees under consideration. Thus, if $\mu_i(t) < \mu_j(t)$ for all $j \neq i$, then for any b with $t_b = t$, i wins b at v . Choose α_i and β_i so that $\min_j(\mu_j(t))$ is a continuous piecewise linear function, $\min_j(\mu_j(t)) = \mu_1(t)$ if $t \in [0, \frac{5}{2}]$, and for $1 < i \leq r$, $\min_j(\mu_j(t)) = \mu_i(t)$ if $t \in [2i - \frac{3}{2}, 2i + \frac{1}{2}]$. Then 1 wins b at v if $t_b = 0, 2$, and for $1 < i \leq r$, i wins b at v if $t_b = 2i$ or $2i - 1$.

Form the matrices $M_i(v)$ with columns equal to the vectors in $D_i(v)$. Now we just need to show that for $1 \leq i \leq r$, $\text{rank}(M_i(v)) = (2n - 3) + 1$. In order to do so, we will reinterpret our matrices in order to utilize previous results about reconstructing trees from subtree weights. Let T' be a tree, and assign to each edge a positive weight $w(e)$. Define the weight of a subforest to be the sum of the weights of the edges contained in the subforest. Let $M'_i(v)$ be the matrix consisting of the first $2n - 3$ rows of $M_i(v)$ and

$$w = \begin{pmatrix} w(e_1) \\ \vdots \\ w(e_{2n-3}) \end{pmatrix},$$

then $M'_i(v)w$ is a column vector with j -th entry equal to the weight of the subforest corresponding to the j -th column of $M'_i(v)$.

$M'_1(v)$ contains column vectors corresponding to the empty subforest as well as the subforests with two leaf edges. A subforest with exactly two leaf edges with degree one vertices u and v is just the path between u and v with weight $d(u, v)$. Therefore, for fixed w , the entries of the column vector $M'_1(v)w$ determine a tree metric δ which has a graph realization T' . By the *Tree-Metric theorem* ([14, 16]) T' is the unique tree metric representation of δ , and w is the unique solution to

$$M'_1(v)x = M'_1(v)w.$$

Therefore, we know $\text{rank}(M'_1(v)) = 2n - 3$. In fact, we have that the column rank of just the columns corresponding to subforests with two leaf edges is equal to $2n - 3$. If we let $(x_1, \dots, x_{2n-3}, y_1, \dots, y_{2n-3})$ be an arbitrary vector in \mathbb{R}^{4n-6} , then each of the columns is contained in the subspace defined by

$$x_1 + \dots + x_n = \frac{2}{n-2}(y_1 + \dots + y_n).$$

The column corresponding to the empty subforest is clearly not contained in this subspace, so its addition increases the column rank by one, which implies $\text{rank}(M_1(v)) = 2n - 2$.

For $i \geq 2$, in order to show that $\text{rank}(M_i(v)) = 2n - 2$, we will first show that we can recover every edge weight if we know the weight of the subforests on $2i$ and $2i - 1$ leaves. To determine the weight of a leaf edge uv with degree three vertex u , choose a subforest with $2i$ leaves that includes all three edges incident to u . Choosing such a subforest is always possible since $2i \geq 3$. Removing uv results in a subforest with $2i - 1$ leaves with corresponding vector also in $D_i(v)$. The difference of the weights of these two subforests determines the weight of the leaf edge.

For an internal edge uv , we construct a subforest that includes the edge uv and the other four edges incident to either u or v . The fact that $2i \geq 4$ ensures that such a subforest exists. Again, omitting uv from this subforest gives us a different subforest on $2i$ leaves, and subtracting, we obtain the weight of uv .

For each edge, we found two subforests that differed by exactly that edge. Subtracting these vectors we obtain every column of the matrix $\begin{pmatrix} I \\ -I \end{pmatrix}$, where I is the $(2n - 3) \times (2n - 3)$ identity matrix. Anything in the column span of these vectors possesses the property that the entry for a_C^k is just the negative of the entry for a_A^k . Therefore, adding any vector without this property to the set, i.e., adding any of the other subforest vectors, increases the rank by one. Thus, $\text{rank}(M_i(v)) = 2n - 2$.

Case 2: r is odd

We construct the vector $v = (v_1, \dots, v_r)$ as in the first case so that $\min_j(\mu_j(t))$ is a continuous piecewise linear function with

$$\min_j(\mu_j(t)) = \begin{cases} \mu_1(t) & : t \in [0, \frac{5}{2}] \\ \mu_2(t) & : t \in [\frac{5}{2}, 4] \\ \mu_3(t) & : t \in [4, \frac{11}{2}] \\ \mu_i(t) & : t \in [2i - \frac{5}{2}, 2i - \frac{1}{2}]. \end{cases}$$

Notice that no i wins v at b if $t_b = 4$, but that if $t_b \neq 4$, then

- 1 wins b at v if $t_b = 0$ or 2
- 2 wins b at v if $t_b = 3$
- 3 wins b at v if $t_b = 5$
- for $3 < i \leq r$, i wins b at v if $t_b = 2i - 2$ or $t_b = 2i - 1$.

Let s_b^3 be the number of internal edges of the subforest induced by b on T_3 . We will perturb the entries of v_3 so that none of the above winning directions is affected but so that if $t_b = 4$, then 3 wins b at v if $s_b^3 = 0$ and 4 wins b at v if $s_b^3 > 0$.

Set the entries of v_3 corresponding to internal edges a_C^j equal to γ and to internal edges a_A^j equal to δ . Let $\tilde{\mu}_3(s, t) = \mu_3(t) + \gamma s + (n - 3 - s)\delta$. Now $l_{3,b}(v_3) = \tilde{\mu}_3(s_b^3, t_b)$ and for $i \neq 3$, $l_{i,b}(v_i) = \mu_i(t_b)$. Let $\delta = \epsilon > 0$ and $\gamma = -\epsilon(n - 3.9)$. Now if $t_b = 4$ and $s_b = 0$, then $l_{3,b}(v_3) = \tilde{\mu}_3(0, 4) = \mu_3(4) + \epsilon(n - 3) = \mu_2(4) + \epsilon(n - 3) = l_{2,b}(v_2) + \epsilon(n - 3) > l_{2,b}(v_2)$, and so 2 wins b at v .

By our choice of γ and δ , as s increases $\tilde{\mu}_3(s, t)$ decreases. Therefore, to show that 3 wins b at v when $t_b = 4$ and $s_b > 0$, it is enough to show that 3 wins b at v when $t_b = 4$ and $s_b = 1$. In that case, we have $l_{3,b}(v_3) = \tilde{\mu}_3(1, 4) = \mu_3(4) + \gamma + \epsilon(n - 4) < \mu_3(4) = \mu_2(4) = l_{2,b}(v_2)$, which implies 3 wins b at v . Now choose $\epsilon > 0$ small enough to ensure that the winning directions for $t_b \neq 4$ are unaffected.

Now we would like to show that for $1 \leq i \leq r$, $\dim_{\mathbb{R}}\langle D_i(v) \rangle_{\mathbb{R}} = 2n - 2$. When $i = 1$, the proof is the same as in the even case. Likewise, for $i > 3$, $D_i(v)$ contains vectors corresponding to subforests with $2i - 2$ and $2i - 1$ leaves. Essentially the same arguments from the even case where $D_i(v)$ contained vectors corresponding to subforests with $2i - 1$ and $2i$ leaves show that $\dim_{\mathbb{R}}\langle D_i(v) \rangle_{\mathbb{R}} = 2n - 2$. Therefore, we just need to establish the rank of $M_2(v)$ and $M_3(v)$.

$M_2'(v)$ contains vectors corresponding to every 3-leaf subtree of T_2 . Just as before, given the weight vector w , $M_2'(v)w$ encodes the weight of every 3-leaf subtree. The main theorem in [14] implies that these weights uniquely determine w , so $\text{rank}(M_2'(v)) = 2n - 3$. Every one of these vectors is contained in the hyperplane defined by

$$x_1 + \dots + x_n = \frac{3}{n-3}(y_1 + \dots + y_n).$$

Since any binary tree with four or more leaves contains at least two cherries, T_3 has a 4-leaf subforest with no internal edges, and 2 wins v at the coordinate corresponding to this subforest. The vector corresponding to this subforest is not contained in the hyperplane above, so $\text{rank}(M_2(v)) = 2n - 2$.

Just as before, for each edge e of T_3 we use two subforests from the set of 4 and 5-leaf subforests that differ only by e to determine $w(e)$. To carryout this procedure, we never require a subforest with no internal edges. This would only be the case if when isolating an internal edge e with endpoints u and v , every other edge adjacent to u and v was a leaf edge. However, that would imply that $n = 4$, which is contradiction. Again, the columns of the matrix $\begin{pmatrix} I \\ -I \end{pmatrix}$ are in the column span and adding any of the original subforest vectors gives us $\text{rank}(M_3(v)) = 2n - 2$. \square

6. PHYLOGENETIC INVARIANTS

As outlined, the fact that $D(3) = 6$ means that to show the generic identifiability of the tree parameters for the 3-tree Jukes-Cantor mixture model, we only need to separate all mixtures with $n = 6$ leaves. While this is at least finite, there are still 19,698,048,370 pairs of 6-leaf triplets, making it infeasible to try to list all pairs and separate them directly. For some pairs $S, T \in \mathcal{T}_{[6],3}$, there is a disentangling set K such that $|K| < 6$, so our strategy is to first generate all pairs of 5-leaf triplets and wherever possible show the mutual non-containment of their corresponding varieties. By application of Theorem 5.3, we actually only need to show the two varieties are not identical, and to do so we will need to find a phylogenetic invariant that holds

for one mixture and not the other. Once complete, we will have a short list of 5-leaf triplet pairs for which we are unable to show that their varieties are not identical. From this list, we arrive at a much smaller list of pairs of 6-leaf triplets which we need to separate using invariants. For these 6-leaf triplets, we use linear invariants and higher order invariants to separate all of the pairs.

This final step in the proof of Theorem 1.1 is highly computational. The steps are all completely contained in the three Maple ([1]) worksheets

`LinearInvariants_5Leaf.mw`, `LinearInvariants_6Leaf.mw`,
`Higher_Degree_Invariants.mw`

which are located at the website:

<http://www4.ncsu.edu/~smsulli2/Pubs/ThreeTreesWebsite/threetrees.html>

We outline the methods for finding separating linear invariants and higher-order invariants in the next two subsections.

6.1. Linear Invariants. Our first step will be to distinguish 3-tree mixture models with 5 leaves from one another by finding linear invariants that hold for one mixture but not the other. A few observations will help us reduce the dimension of the ambient space of the varieties. For our purposes, it is unnecessary to calculate invariants that hold for all mixtures. For the Jukes-Cantor model, because $\{C, G, T\}$ are in the same class, every model will have linear invariants arising from permutations on this set. For example, any five leaf mixture will have the same parameterization on the coordinates in the set

$$\{qCCCGT, qGGGCT, qTTTGC, qCCCTG, qGGGTC, qTTTCG\}.$$

Therefore, the difference of any two of these elements is an invariant for every 3-tree mixture model with 5 leaves. We will only consider the lexicographically first element as a representative of each such set. By doing this, and removing coordinates that are always zero, we can perform our calculations in \mathbb{C}^{51} instead of \mathbb{C}^{1024} . (In the provided maple file, we also exclude the coordinate $qAAAAA$, which is not involved in any linear invariant).

Applying an element of S_5 to the leaf labels of all of the trees in a mixture model merely permutes the coordinates of our parameterization. As a result, once we have determined that the models corresponding to a 5-leaf triplet pair do not in fact have the same variety, then by applying the elements of S_5 to all six trees in the pair, we generate new 5-leaf triplet pairs that do not have the same variety. To illustrate, although there are 680 different identical 5-leaf triplet pairs $((T, T) \in \mathcal{T}_{[n],r} \times \mathcal{T}_{[n],r})$, all of these pairs can be generated by applying an element of S_5 to one of only 28 such pairs.

The following lemma allows us to compute the linear invariant space of every 3-tree mixture model using the built-in Maple function `IntersectionBasis`.

Lemma 6.1. *A linear polynomial f is an invariant of $V_{T_1} * V_{T_2} * V_{T_3} \iff f$ is an invariant of V_{T_1}, V_{T_2} , and V_{T_3} .*

In the provided code, we compare all 5-leaf triplet pairs, and up to the action of S_5 , there are 36 pairs with the same linear invariant space. As mentioned, this list consists of the 28 pairs of identical triplets as well as 8 additional pairs containing two distinct triplets.

If there exists $S, T \in \mathcal{T}_{[6],3}$ such that $V_{T_1} * V_{T_2} * V_{T_3} = V_{S_1} * V_{S_2} * V_{S_3}$, then it must be the case that for any five element subset K , $(S|_K, T|_K)$ (or some permutation of the leaves thereof) is one of the 5-leaf triplet pairs in our list. Therefore, the only 6-leaf triplet pairs that are candidates for inseparability are those generated by attaching an additional edge to each of the 6 trees in an inseparable 5-leaf triplet pair. Since there are 36 pairs, and each 5-leaf tree has 7 edges, the number of 6-leaf triplets we must consider is less than $(36)(7^6) = 4,235,364$. This is far fewer than would be expected had we proceeded directly to the 6-leaf case.




FIGURE 5. A 6-leaf triplet pair that is not separated by linear invariants.

Just as in the 5-leaf case, nondefectiveness of all of the involved varieties is ensured by Theorem 5.3, so it suffices to show that $V_{T_1} * V_{T_2} * V_{T_3} \neq V_{S_1} * V_{S_2} * V_{S_3}$. This fact is reflected in our computation of the set `AllSixLeafPairs`, which contains the eighty-five 6-leaf triplet pairs (up to the action of S_6) that are not separated by any linear invariant. For these eighty-five pairs of 6-leaf triplets we must find higher degree invariants that separate them.

6.2. Invariants of higher degree. If we wish to determine identifiability, we must broaden our search to invariants of higher degree. Unfortunately, we glean little from studying V_{T_1} , V_{T_2} , and V_{T_3} individually, as there is no reason to expect that a non-linear invariant that holds for all three holds for their join.

After removing trivial coordinates and linear invariants that hold for all 6-leaf trees we can perform our calculations for 6-leaf trees in 186-dimensional image space. When searching for linear invariants, we will also disregard any coordinate in which the character A appears in the index so that the bulk of our computations are done in a 31-dimensional image space. If we let T be the six leaf triplet described in Figure 5 (in the worksheet, these trees are labelled T_{16} , T_{19} and T_{63}), there are 61 linearly independent elements of $I(V_{T_1} * V_{T_2} * V_{T_3})_1$ which define a 125-dimensional linear subspace containing $V_{T_1} * V_{T_2} * V_{T_3}$. Since $\dim(V_{T_1} * V_{T_2} * V_{T_3}) = 29$, at least for this particular triplet, higher-degree invariants must exist. In order to find these, we will use the explicit parameterization of the variety in Fourier coordinates. Here, for example, is the parameterization of T on just the first three coordinates. As mentioned in Section 4, we make the expressions more compact by letting $a_C^{i|[6]\setminus i} = a_i$, $b_C^{i|[6]\setminus i} = b_i$, $c_C^{i|[6]\setminus i} = c_i$, $a_C^{15|2346} = a_7$, $a_C^{15|2346} = a_8$, $a_C^{135|246} = a_9$, $b_C^{12|3456} = b_7$, $b_C^{15|2346} = b_8$, $b_C^{126|345} = b_9$, $c_C^{12|3456} = c_7$, $c_C^{34|1256} = c_8$, and $c_C^{125|346} = c_9$.

$$\begin{aligned} q_{AAAAACC} &= \pi_1 a_5 a_6 a_7 a_8 a_9 + \pi_2 b_5 b_6 b_8 b_9 + \pi_3 c_5 c_6 c_9 \\ q_{AAACAC} &= \pi_1 a_4 a_6 + \pi_2 b_4 b_6 b_9 + \pi_3 c_4 c_6 c_8 \\ q_{AAACCA} &= \pi_1 a_4 a_5 a_7 a_8 a_9 + \pi_2 b_4 b_5 b_8 + \pi_3 c_4 c_5 c_8 c_9 \end{aligned}$$

Theoretically, we should be able use these equations to construct an ideal in

$$\mathbb{C}[y_1, \dots, y_{186}, a_1, \dots, a_9, b_1, \dots, b_9, c_1, \dots, c_9, \pi_1, \pi_2],$$

and eliminate to obtain a Gröbner basis for $I(V_{T_1} * V_{T_2} * V_{T_3})$ that includes non-linear invariants. The number of variables and equations involved in the computation make it infeasible.

Instead, we will apply heuristic methods to reduce the number of variables in the ideal before we attempt elimination. Our strategy will be to find linear invariants that hold for V_{T_2} and V_{T_3} , but not for V_{T_1} . The resulting equations will not evaluate to zero on all of $V_{T_1} * V_{T_2} * V_{T_3}$, but will not involve the parameters from T_2 or T_3 at all. To illustrate, we revisit the trees in Figure 5, and look at a few more coordinates.

$$\begin{aligned} qCCCAAC &= \pi_1 a_1 a_2 a_3 a_6 a_7 a_8 + \pi_2 b_1 b_2 b_3 b_6 b_8 b_9 + \pi_3 c_1 c_2 c_3 c_6 c_8 \\ qCCGAAG &= \pi_1 a_1 a_2 a_3 a_6 a_7 a_8 a_9 + \pi_2 b_1 b_2 b_3 b_6 b_8 b_9 + \pi_3 c_1 c_2 c_3 c_6 c_8 \\ qCACCGT &= \pi_1 a_1 a_3 a_4 a_5 a_6 a_7 a_8 a_9 + \pi_2 b_1 b_3 b_4 b_5 b_6 b_7 b_8 b_9 + \pi_3 c_1 c_3 c_4 c_5 c_6 c_7 c_9 \\ qCAGGTG &= \pi_1 a_1 a_3 a_4 a_5 a_6 a_7 + \pi_2 b_1 b_3 b_4 b_5 b_6 b_7 b_8 b_9 + \pi_3 c_1 c_3 c_4 c_5 c_6 c_7 c_9, \end{aligned}$$

Now it is easy to spot linear invariants for V_{T_2} and V_{T_3} , and subtracting we obtain

$$\begin{aligned} qCCCAAC - qCCGAAG &= \pi_1 a_1 a_2 a_3 a_6 a_7 a_8 - \pi_1 a_1 a_2 a_3 a_6 a_7 a_8 a_9 \\ qCACCGT - qCAGGTG &= \pi_1 a_1 a_3 a_4 a_5 a_6 a_7 a_8 a_9 - \pi_1 a_1 a_3 a_4 a_5 a_6 a_7. \end{aligned}$$

In this particular case, $\dim(I(V_{T_1} * V_{T_2} * V_{T_3})_1) - \dim(I(V_{T_2} * V_{T_3})_1) = 20$, so there are twenty linearly independent relations only involving the parameters from T_1 . We introduce new variables for the image space and use these relations to construct the ideal,

$$J = \langle y_1 - (\pi_1 a_1 a_2 a_3 a_6 a_7 a_8 - \pi_1 a_1 a_2 a_3 a_6 a_7 a_8 a_9), y_2 - (\pi_1 a_1 a_3 a_4 a_5 a_6 a_7 a_8 a_9 - \pi_1 a_1 a_3 a_4 a_5 a_6 a_7), \dots \rangle,$$

where $J \subseteq \mathbb{C}[y_1, \dots, y_{20}, a_1, \dots, a_9, \pi_1]$. With fewer parameters, we can now compute a Gröbner basis for $J \cap \mathbb{C}[y_1, \dots, y_{20}]$ using elimination in Macaulay2 [10]. The basis gives us a set of relations in the y_i variables, which we translate back into our original coordinates. For this particular triplet, we find

$$(qCCCAAC - qCCGAAG)(qCACCGT - qCAGGTG) = (qCCAGT - qCCGATC)(qCACCAC - qCAGGAC).$$

Finally, to separate the triplet pair from Figure 5, we substitute the parameterization of S into this relation and confirm that it does not evaluate to zero.

This technique allows us to find an invariant for one mixture that does not hold for the other for all of the triplet pairs contained in `AllSixLeafPairs`. As outlined, the existence of these invariants is sufficient to establish the generic identifiability of the tree parameters of the 3-class Jukes-Cantor mixture model. Among the supplementary materials is the worksheet `HigherDegreeInvariants` which lists an invariant separating each triplet pair and provides code to quickly generate the coordinate functions for verification.

ACKNOWLEDGMENTS

Colby Long was partially supported by the US National Science Foundation (DMS 0954865). Seth Sullivant was partially supported by the David and Lucille Packard Foundation and the US National Science Foundation (DMS 0954865).

REFERENCES

- [1] Maple 14. Maplesoft, a division of Waterloo Maple Inc. Waterloo, Ontario.
- [2] E. Allman, C. Ané, and J.A. Rhodes. Identifiability of a markovian model of a molecular evolution with gamma-distributed rates. *Adv. Appl. Prob.*, 40:229–249, 2008.
- [3] E. Allman and J.A. Rhodes. Identifying evolutionary trees and substitution parameters for the general markov model with invariable sites. *Math. Biosci.*, 211(1):18–33, 2008.
- [4] E.S. Allman, S. Petrovic, J.A. Rhodes, and S. Sullivant. Identifiability of 2-tree mixtures for group-based models. *IEEE/ACM Trans Comput Biol Bioinformatics*, 8(3):710–722, 2011.
- [5] E.S. Allman and J.A. Rhodes. The identifiability of tree topology for phylogenetic models, including covarion and mixture models. *J. Comp. Biol.*, 13(5):1101–1113, 2006.

- [6] E.S. Allman, J.A. Rhodes, and S. Sullivant. When do phylogenetic mixture models mimic other phylogenetic models? *Syst. Biol.*, 61(6):1049–1059, 2012.
- [7] J.T. Chang. Full reconstruction of markov models on evolutionary trees: identifiability and consistency. *Math. Biosci.*, 137(1):51–73, 1996.
- [8] J. Draisma. A tropical approach to secant dimensions. *J. Pure Appl. Algebra*, 212(2):349–363, 2008.
- [9] S.N. Evans and T.P. Speed. Invariants of some probability models used in phylogenetic inference. *Ann. Statist.*, 21(1):355–377, 1993.
- [10] D.R. Grayson and M.E. Stillman. Macaulay2, a software system for research in algebraic geometry. Available at <http://www.math.uiuc.edu/Macaulay2/>, 2002.
- [11] P.J. Humphries. *Combinatorial aspects of leaf-labelled trees*. PhD thesis, University of Canterbury, 2008.
- [12] F.A. Matsen, E. Mossel, and M. Steel. Mixed-up trees: the structure of phylogenetic mixtures. *Bull. Math Biol.*, 70(4):1115–1139, 2008.
- [13] F.A. Matsen and M. Steel. Phylogenetic mixtures on a single tree can mimic a tree of another topology. *Syst. Biol.*, 56(5):767–775, 2007.
- [14] L. Pachter and D. Speyer. Reconstructing trees from subtree weights. *Applied Mathematical Letters*, 17:615–621, 2004.
- [15] J.A. Rhodes and S. Sullivant. Identifiability of large phylogenetic mixtures. *Bull. Math Biol.*, 74(1):212–231, 2012.
- [16] C. Semple and M. Steel. *Phylogenetics*. Oxford University Press, Oxford, 2003.
- [17] M.A. Steel and Y.X. Fu. Classifying and counting linear phylogenetic invariants for the jukes-cantor mixture model. *J. Comp. Biol.*, 2(1):39–47, 1995.
- [18] B. Sturmfels and S. Sullivant. Toric ideals of phylogenetic invariants. *J. Comp. Biol.*, 12(2):204–228, 2005.
- [19] S. Sullivant. The disentangling number for phylogenetic mixtures. *SIAM J. Discrete Math*, 26(2):856–859, 2012.
- [20] L. Székely, P.L. Erdős, M.A. Steel, and D. Penny. A fourier inversion formula for evolutionary trees. *Applied Mathematics Letters*, 6(2):13–17, 1993.

E-mail address: celong2@ncsu.edu

E-mail address: smsulli2@ncsu.edu

DEPARTMENT OF MATHEMATICS, BOX 8205, NORTH CAROLINA STATE UNIVERSITY, RALEIGH, NC, 27695-8205, USA